A Bayesian Approach to Microwave Precipitation Profile Retrieval

K. FRANKLIN EVANS

Department of Atmospheric Science, Colorado State University, Fort Collins, Colorado

JOSEPH TURK

Department of Electrical Engineering, Colorado State University, Fort Collins, Colorado

TAKMENG WONG AND GRAEME L. STEPHENS

Department of Atmospheric Science, Colorado State University, Fort Collins, Colorado

(Manuscript received 9 December 1993, in final form 1 April 1994)

ABSTRACT

A multichannel passive microwave precipitation retrieval algorithm is developed. Bayes theorem is used to combine statistical information from numerical cloud models with forward radiative transfer modeling. A multivariate lognormal prior probability distribution contains the covariance information about hydrometeor distributions that resolves the nonuniqueness inherent in the inversion process. Hydrometeor profiles are retrieved by maximizing the posterior probability density for each vector of observations. The hydrometeor profile retrieval method is tested with data from the Advanced Microwave Precipitation Radiometer (10, 19, 37, and 85 GHz) of convection over ocean and land in Florida. The CP-2 multiparameter radar data are used to verify the retrieved profiles. The results show that the method can retrieve approximate hydrometeor profiles, with larger errors over land than water. There is considerably greater accuracy in the retrieval of integrated hydrometeor contents than of profiles. Many of the retrieval errors are traced to problems with the cloud model microphysical information, and future improvements to the algorithm are suggested.

1. Introduction

The promise of satellite-based passive microwave remote sensing of precipitation has been that accurate rainfall retrieval would be possible because microwave radiation penetrates clouds and is a direct measure of the underlying rain rate. As the field of microwave precipitation retrieval has been explored, evidence has accumulated indicating that the situation is more complex. Empirical statistical methods that relate observed microwave brightness temperatures to surface rain rate (e.g., measured with radar) necessarily avoid considering the complexities of the interaction of microwave radiation with real precipitating systems. These methods, while using a perfectly valid approach, have been hampered by a lack of representative and accurate ground truth. The microwave precipitation retrieval methods that are based on radiative transfer calculations must have a model of the vertical structure of precipitation, including the relevant hydrometeor species. The prototypical simple physical algorithm is the 19-GHz emission method of Wilheit et al. (1977). In its earliest form, this method assumed a uniform rain layer with a Marshall–Palmer distribution below the freezing level with no ice particles and a prescribed liquid water cloud. These assumptions allow a simple relationship between the 19-GHz brightness temperature and the rain rate to be derived from a radiative transfer model. Many other microwave precipitation algorithms that assume very simple precipitation structures have been developed (e.g., Weinman and Guetter 1977; Olson 1989; Liu and Curry 1992). A common characteristic of most of these simple physical models is the use of one observable, perhaps a combination of several channels, that is related to surface rain rate by assuming some type of hypothetical hydrometeor profile. Typically, there are a few homogeneous layers and the mass contents of the hydrometeors are slaved to the surface rain rate.

There have been several studies that have led to a deeper understanding of how microwave radiation interacts with realistic profiles of hydrometeors. Fulton and Heymsfield (1991) qualitatively compared hydrometeor information inferred from multiparameter radar data of intense convection with brightness temperatures at 18, 37, 92, and 183 GHz. Their results suggest that even the lowest frequency (18 GHz) is sig-

Corresponding author address: Frank Evans, PAOS, Campus Box 311, University of Colorado, Boulder, CO 80309.

nificantly obscured by the large ice mass in deep convection. Yeh et al. (1990) and Vivekanandan et al. (1990) used radiative transfer models to simulate microwave brightness temperatures of radar-derived hydrometeor profiles and compared the results with brightness temperatures observed from aircraft. The use of microwave radiative transfer simulations of hydrometeor profiles derived from cloud models has led to a number of useful insights. Adler et al. (1991) used a three-dimensional simulation of a tropical oceanic squall line to investigate brightness temperature–rain rate relations at frequencies from 10 to 85 GHz. They showed a large scatter in the relations for 19 GHz and above due to variations in ice content and cloud liquid water, which vary systematically with the nature (convective or stratiform) of the precipitation. Mugnai et al. (1990), and more recently Smith et al. (1992) and Mugnai et al. (1993), used a bulk microphysics cloud model and radiative transfer model to simulate frequencies from 6 to 128 GHz. Their use of weighting functions allowed a detailed examination of the effects of the vertical distribution of various types of hydrometeors on the upwelling microwave radiation. They found that most frequencies respond primarily to fluctuations in graupel mass.

These modeling studies have shown that passive microwave measurements of precipitating clouds are sensitive to many aspects of the vertical distribution of various hydrometeor species and not solely to the surface rain rate. Since the microwave radiation is sensing the profile of hydrometeors and a number of frequencies are potentially available, it should be profitable to retrieve hydrometeor profiles rather than only the surface rain rate. Profile retrieving algorithms are designed to use all the available channels rather than just the one or two that the simple models use. The vertical distribution of hydrometeors is important in its own right because it can be related to the vertical profile of latent heat release that governs the coupling to dynamical forcing (Tao et al. 1990). Profiling algorithms will work over land as well as water, although with reduced accuracy. Of course, simple physical models have their uses when speed or simplicity are the driving concern, but profiling algorithms potentially make the best use of all the available data.

The first profile retrieving method was that of Kummerow (Kummerow et al. 1989; Kummerow et al. 1991). This method used an iterative scheme that matches the observed brightness temperatures with those simulated from a relatively small number of somewhat ad hoc specified profiles. This method is the furthest advanced of the profiling algorithms and has been used successfully in a variety of situations (Kummerow and Giglio 1994a, 1994b). Obviously one cannot expect to retrieve highly accurate detailed profiles from a handful of microwave channels. In fact, since there are typically multiple distinct profiles that can satisfy a small set of observations, other information about the vertical distribution of hydrometeors is needed to constrain the retrieval. Mugnai et al. (1993) have developed a profile retrieval method that uses a database of cloud model–derived profiles and their corresponding simulated brightness temperatures. Profiles with brightness temperatures close to the observed ones are selected and combined using a weighted average with the weights adjusted so as to match the observations. A closely related algorithm that uses the same cloud–radiation database (Marzano et al. 1994) has been tested with the same land dataset as in the validation experiment in section 4.

The profile-retrieving algorithm developed here also uses cloud model information, though in a significantly different way. The cloud model hydrometeor profiles are used statistically to generate a probability density function that represents the prior information about microphysical structure. Forward radiative transfer modeling of brightness temperatures is used to define the conditional probability of the observations given a particular hydrometeor profile. Bayes theorem combines the prior distribution and the conditional distribution to find the posterior probability density function of the profile given the observation. In this work, the most likely profile that matches the observations is then found from the posterior probability function. A more complete description of the method can be found in Evans and Stephens (1993). Compared with the cloud–radiation database approach, the Bayesian method has the advantage of having a more formal way of introducing cloud microphysical information, which aids in analyzing the behavior of the algorithm. The Bayesian approach can generalize beyond the particular cloud model profiles to analogous situations and picks the a priori likeliest profiles rather than using all those that are close to matching the observations. The Bayesian retrieval method does not match the observations exactly, so the results are less affected by errors in the radiative transfer modeling. On the other hand, the Bayesian method is dependent on choosing appropriate forms for the probability distributions and on the method by which the profiles are retrieved from the posterior distribution. The Bayesian approach is not guaranteed to produce reasonable-looking smooth profiles as is the cloud–radiation database method. Both methods depend critically on having correct microphysical information from cloud models.

A detailed description of the Bayesian precipitation retrieval algorithm is given in the next section. A theoretical test of the method using brightness temperatures simulated from cloud model output is done to determine the complexity required of the precipitation profile. A limited validation experiment with data from the Convection and Precipitation/Electrification (CaPE) experiment is carried out. Brightness temperatures from the Advanced Microwave Precipitation

Radiometer (AMPR) are used to retrieve hydrometeor profiles, which are compared with CP-2 multiparameter radar data. The reasons for the successes and failures of the algorithm are investigated and areas for improvement of the method are discussed.

2. Bayesian precipitation retrieval method

a. *The Bayesian framework*

This retrieval method uses Bayes theorem to combine forward radiative transfer modeling with statistical information from numerical cloud model output. For use here, Bayes theorem may be stated as

$$f_{\Theta|X}(\boldsymbol{\theta}|\mathbf{x}) \propto f_{X|\Theta}(\mathbf{x}|\boldsymbol{\theta}) f_{\Theta}(\boldsymbol{\theta}), \tag{1}$$

where $\boldsymbol{\theta}$ represents the atmospheric state vector (i.e., hydrometeor profile) and $\mathbf{x}$ represents the vector of observations (e.g., microwave T_b). The prior probability distribution of atmospheric parameters is $f_{\Theta}(\boldsymbol{\theta})$, $f_{X|\Theta}(\mathbf{x}|\boldsymbol{\theta})$ is the conditional probability distribution of an observation given an atmospheric state, and $f_{\Theta|X}(\boldsymbol{\theta}|\mathbf{x})$ is the posterior probability distribution function of an atmospheric state given the observation. The normalization of the posterior distribution has been left out, hence the proportionality rather than equality. The conditional distribution, which is the probability of an observation given an atmospheric state, is closely related to the forward problem, that is, computing observables from a given atmosphere using radiative transfer. The prior distribution is the vehicle for extra information about the atmospheric state that will reduce the ill-conditioned nature of the inversion. Since a single profile is usually desired for a retrieval, the maximum or the expected value of the posterior probability distribution can be used.

There are several advantages to using Bayes theorem in a retrieval method:

1) Bayes theorem provides an elegant mathematical framework for adding prior information about precipitating systems to improve the accuracy of retrieval.

2) Different types of observational data (e.g., microwave, radar, infrared) can be incorporated easily because only the relatively well understood forward problem needs to be reformulated for each type of data.

3) Estimates of the uncertainties of the parameters, as well as the values, can potentially be computed, because the complete posterior probability density function is available.

4) If the probability distributions are formed correctly, then using Bayes theorem offers the choice of finding either the most likely atmospheric state or the unbiased mean state, as appropriate.

While Bayes theorem provides an overall framework, the form of the prior and conditional probability density functions (pdf's) must still be specified. In the Bayesian framework, probability distributions do not represent the frequency of occurrence but rather our degree of belief or understanding of a system. Thus, the functional forms of the pdf's are to some extent arbitrary. There is a trade-off between simplicity and realism in the functional forms of the Bayesian pdf's, with the goal to capture the necessary behavior with as few parameters as possible. The conditional probability distribution is constructed simply by assuming that the observation vector $\mathbf{x}$ is normally distributed around the simulated observation vector calculated from the forward model $\mathbf{g}(\boldsymbol{\theta})$:

$$f_{X|\Theta}(\mathbf{x}|\boldsymbol{\theta}) = N[\mathbf{x} - \mathbf{g}(\boldsymbol{\theta}), \sigma]. \tag{2}$$

The observation uncertainties (σ_m for each element m of the observation vector) can be thought of as a combination of the actual observation errors and the errors in the forward model. Independence is probably a good assumption for the observation errors but may not be so appropriate for the forward modeling errors. The maximization of the posterior probability function, by also tending to maximize the conditional probability, assures that the simulated observations are close to the actual observations. It does not, however, fit the retrieval to the noise in the observations or to the modeling error if the σ's are chosen appropriately. The σ's can be thought of as controlling the trade-off between fitting more to the data versus being closer to the prior distribution.

Use of a prior probability distribution of atmospheric parameters is the key improvement of the Bayesian retrieval algorithm over previous precipitation profile retrieval methods. The prior information about the atmosphere helps the retrieval choose the best precipitation structure of the many that satisfy the observations. For this work, the prior probability distribution is assumed to be a multivariate lognormal distribution. The hydrometeor species most important for microwave observations are known to have approximately lognormal distributions, as studies of radar data from GARP (Global Atmospheric Research Program) Atlantic Tropical Experiment (GATE) have shown (Kedem et al. 1990). The multivariate aspect of the prior probability distribution is crucial because this is how information about the correlations between hydrometeors at various levels is introduced. The form of the prior distribution is

$$f_{\Theta}(\boldsymbol{\theta}) = \frac{1}{[(2\pi)^n \det\mathbf{C}]^{1/2} \prod_{i=1}^{n} \theta_i} \times \exp\left[-\frac{1}{2}(\mathbf{u} - \bar{\mathbf{u}})^{\mathrm{T}}\mathbf{C}^{-1}(\mathbf{u} - \bar{\mathbf{u}})\right], \tag{3}$$

where $\mathbf{u} = \ln\boldsymbol{\theta}$ is the log of the precipitation parameter vector, $\bar{\mathbf{u}}$ is the first log moment (mean vector) of the

atmospheric vector, and **C** is the central second log moment (covariance matrix) of the atmospheric vector. The constants $\bar{\mathbf{u}}$ and **C** are measured statistically from the numerical cloud model output. This functional form is the simplest way to introduce correlations between variables. However, it cannot represent nonlinear relations between variables, bimodalities, etc. One reason to use a simple function like this, rather than a higher-order function, is that a more general function having more parameters requires substantially more data for a reliable fit.

Since the Bayesian precipitation retrieval method uses a radiative transfer model to simulate observations, it is necessary to incorporate an underlying model of a precipitating atmosphere. The structure of the precipitating atmosphere may be chosen with a high degree of flexibility as to the number and thicknesses of the layers and what parameters are variable in the layers. There is an important distinction between the total set of atmospheric parameters that are used in the forward radiative transfer calculation and the subset of those parameters that are free to be adjusted in the retrieval process. A layer is specified by its thickness, interface temperature, water vapor mass content, cloud liquid water mass content, and the parameters that describe the precipitating hydrometeor distributions. The rain and up to two species of ice particles are assumed to have exponential size distributions. The mass content and the average particle diameter specify the two parameters of the distribution. In addition, the bulk density of the ice particles is a parameter, so that two different ice species may be simulated (e.g., graupel and snow). The particles are assumed to be spherical since many precipitation-size hydrometeors are roughly spherical or randomly oriented, and the errors introduced by this assumption are small compared to other errors.

Besides the parameters that define each layer, there are also parameters for the surface temperature and emissivity. Because of the lack of prior information about surface parameters, such as ocean surface wind speed and soil moisture, the ocean is assumed to be flat and a frequency-independent emissivity is used for land. A precipitation structure is set up for a retrieval by fixing the number of layers and choosing the atmospheric and surface parameters that will be allowed to vary. This defines the parameters making up the atmospheric state vector. Varying all the parameters in a precipitation structure with many layers would lead to an impossibly large number of free parameters, so a judicious choice has to be made. Section 3 tests a number of precipitation structures to determine the best trade-off between the convenience of a few parameters versus the more accurate representation of many parameters.

Numerical cloud models are the only feasible source of the detailed information on the vertical distribution of hydrometeors that is needed to make a suitable prior distribution. Radar data, especially multiparameter data, are also a source of detailed information, but it cannot provide cloud liquid water amounts, size distribution information for ice particles, or distinguish very well between types of ice particles. The cloud models used here have one parameter distributions for the precipitation species (i.e., mixing ratio is predicted), so mass content is the only hydrometeor parameter that is variable in the retrieval. The precipitation structure parameters for each layer are computed from the cloud model output by averaging over the appropriate cloud model levels. The mean vector and covariance matrix of parameters, which determines the multivariate lognormal prior distribution, are calculated from the set of cloud model profiles. All hydrometeor mass contents are clipped at 10^{-4} g m^{-3} to prevent the zero mass points from contaminating the lognormal distribution parameter estimation too much. A vertically integrated rain mass cutoff is used to avoid nonraining pixels that do not belong to the underlying lognormal distribution. Distributions of parameters not obtainable from cloud model output can be specified with a mean and standard deviation.

b. *Forward radiative transfer modeling*

The forward radiative transfer model is the core of the conditional probability distribution part of the precipitation retrieval method. The forward model calculates the gaseous absorption, particle scattering, and radiative transfer from the precipitation structure for each channel in the observation vector. Any number of channels, each with specified frequency, angle, polarization, and observation width (σ_m), may make up the passive microwave observation vector.

The microwave absorption due to oxygen, water vapor, and cloud water are computed using Liebe's MPM92 model (Liebe 1989; Liebe et al. 1993) and are stored as coefficients before proceeding with the retrieval. The absorption coefficient for water vapor is obtained by computing the extinction at 0% and 100% relative humidity for each layer and frequency. The forward model then uses these coefficients by computing the average relative humidity of a layer.

The microwave scattering properties of the exponential distributions of precipitating hydrometeors are found by interpolating from tables of precomputed Mie scattering results. An interpolation approach is necessary in view of the many forward model computations that must be carried out for each retrieval. All of the tables are two dimensional, and the scattering quantities (extinction, single-scattering albedo, and asymmetry parameter) are bilinearly interpolated. For rain, the two dimensions are average diameter and temperature, and for ice they are average diameter and particle density. The Mie tables are computed using

an adaptive scheme that varies the number of grids of average diameter to ensure that the interpolation from the tables will be within a desired tolerance (1%–2% used here). The index of refraction of water is computed following Ray (1972) and for ice is interpolated from the tables in Warren (1984). For ice particles with a density lower than solid ice, the index of refraction is reduced according to the ice volume fraction using the Lorentz–Lorenz (Debye) mixing rule. The Mie calculations are performed by integrating over 100 steps up to a maximum diameter of 10 times the average.

If the mass content of a hydrometeor species is a variable in the precipitation structure, but the average size parameter is not, then either the average size of the exponential distribution is fixed at a specified value or it may be related to the mass content using the Rutledge and Hobbs (1984) parameterization. This scheme fixes the intercept N_0 of the exponential distribution and computes the average diameter parameter from the mass content and the particle density. For rain, N_0 is 8000 mm^{-1} m^{-3} and 4000 mm^{-1} m^{-3} for snow or graupel.

The radiative transfer model used for simulating passive microwave measurements is an Eddington-type plane-parallel two-stream model (e.g., Weinman and Davies 1978). The two-stream model, while less accurate than a multistream model or a model that takes into account the true three-dimensional structure, is necessary in view of the computational burden of the many forward calculations the retrieval method requires. Kummerow (1993) has carried out a comparison between the Eddington approximation and a plane-parallel multistream (discrete ordinate) model, showing that the differences in microwave brightness temperatures are less than 3 K for realistic precipitating cloud profiles. The plane-parallel assumption is likely to be a more significant problem, but we leave for the future the development of parameterizations for dealing with this complicated issue. The Rayleigh–Jeans approximation is used so the transfer calculations are done in terms of brightness temperatures.

When modeling low-density ice particles at 85 GHz, there are large errors in the Eddington approximation. The reason for these errors is the high asymmetry parameter computed by Mie theory for the low index of refraction obtained from the mixing rule. The asymmetry parameter is greater than 0.85 for low densities (0.1 g cm^{-3}) and mean diameters larger than 1.00 mm. For near-nadir angles, the Eddington approximation gives poor results for an asymmetry parameter greater than around 0.6. The Eddington comparison in Kummerow (1993) considered only Marshall–Palmer distributions of solid ice and thus did not have this problem. The standard method for dealing with high asymmetry parameters is to use the delta-Eddington method (Joseph et al. 1976). This modification represents the highly peaked phase function by a delta function and the normal Eddington phase function with a lower asymmetry parameter. The resulting scaled optical properties can be stored in the tables, and the regular Eddington radiative transfer model used unmodified.

The boundary conditions for the radiative transfer model are the 2.7-K blackbody cosmic background radiation from above and surface emission and reflection from below. The surface type is assumed to be known as either land or water. The land surface is modeled as a Lambertian reflector with a frequency-independent emissivity. The water surface is modeled as a flat Fresnel reflector whose index of refraction depends on frequency and the climatological temperature. The only source of polarization in this radiative transfer model is from a water surface where the emissivity depends on the polarization. For simulating observations from the AMPR instrument, which mixes the polarizations as it scans, the final brightness temperature is a linear combination of V (vertical) and H (horizontal) polarizations according to the observation angle.

c. *Retrieval of hydrometeor profiles*

From the prior probability distribution and the forward model, it is simple to compute the posterior distribution from (1) for any given atmospheric state vector. What is needed, however, is to find the single atmospheric state that represents the best retrieved profile from the high-dimensional posterior function. For single-pixel retrievals that are to be validated against another data source, we think the most appropriate approach is to find the "most likely" profile by maximizing the posterior probability function. On the other hand, for remote sensing the mean precipitation profile over a given area and time, the best approach is to integrate the posterior density function to find the moments that specify the distribution. One could use the expected value for single-pixel retrievals as well, but the simulated observations from these retrievals can be quite far from the actual observations.

The appropriate method for finding the mean profile over a set of observations is to estimate the parameters of an assumed lognormal distribution of precipitation. The correct procedure for doing this is to integrate over the posterior distribution, to compute the first and second log moments for each observation, and average these moments over the set of observations, thereby improving the estimate of the distribution parameters. This method is not biased to lower values of rain rate—for example, by missing the rare large rain events—as is simple averaging of pixel-by-pixel retrievals. Because of the high dimensionality of the posterior density function, the only feasible way to integrate over it is to use Monte Carlo integration. Because the prior probability distribution does not change over a set of retrievals and is a simple function, it can be absorbed into the sampling of the points; that is, the Monte Carlo

random points (profiles) are chosen according to the multivariate lognormal prior distribution. Although a large number of random points are needed to achieve adequate accuracy, a set of observation pixels may be processed in parallel by evaluating the forward model only once for each random profile. This method was implemented but not used here because single-pixel retrievals are desired.

The approach used for the retrievals shown below is to maximize the posterior probability density function. There is an appeal in finding the maximum of the posterior distribution because then the retrieval is the most likely atmospheric state. There is, however, some arbitrariness in "maximizing probability." The probability of a continuous distribution at a single point is zero; only ranges have nonzero probabilities. Furthermore, the maximum of a probability density function depends upon what space the density function is expressed in. This can be seen by dividing the density function up into discrete bins; which bin has the maximum probability then depends on how the bins are spaced. Here we use logarithmically spaced ranges. It is convenient to define an objective function J to be minimized, which (to within a constant) is the negative of the log of the posterior function. The objective function is then

$$J = (\mathbf{u} - \bar{\mathbf{u}})^{\mathrm{T}}\mathbf{C}^{-1}(\mathbf{u} - \bar{\mathbf{u}}) + \sum_{m=1}^{N_{\mathrm{obs}}} \frac{[x_m - g_m(\theta)]^2}{\sigma_m^2}. \quad (4)$$

A somewhat complex algorithm is used to minimize the objective function in an attempt to deal with the problem of multiple local minima. The prior density function is very smooth and relatively broad. In the usual case of many more variable parameters than observations, the forward objective function has a sharp valley that is presumably zero in a large subspace; that is, there are many precipitation profiles that exactly satisfy the observations. The purpose of the prior distribution is to lift this degeneracy so there will be one lowest point. There can be (and are) local minima that trap traditional optimization techniques. An optimization method that can deal with local minima, called "simulated annealing," was tried and found to be too slow. Since the objective function is basically smooth, a modification of traditional optimization algorithms was found to be effective.

The idea behind the algorithm is to start off a standard optimization routine from a number of places that are likely to be close to the desired global minimum. This is done by precomputing the simulated observations with the forward model for a large set of precipitation profiles (typically 10^6) chosen from a multidimensional lognormal distribution. For each observation vector, those profiles that match the observations to within a certain distance are selected. These are sorted by the value of the prior probability function and several of the highest prior profiles are used to initialize Powell's direction set method of optimization (see Acton 1990).

The purpose of doing the minimization a number of times from different starting points (5–15 used here) is to attempt to make sure that the resulting lowest function value is close to the true global minimum. A theoretical test (similar to the one in the next section) was done to try to determine if this method overcomes the problem of multiple local minima. Optimizations with just a few starting points give retrieved profiles very close to those for 32 initial points. Doing multiple optimizations does mean that the retrieval algorithm takes an excessive amount of computer time. The number of function evaluations made by the optimization method goes as the square of the number of variables. In the CaPE validation test below, an average of 14 000 posterior probability evaluations were made for each pixel. Since the required minimum of five optimization tries nearly always got the same results, a single optimization probably would be adequate, giving 2800 function evaluations. Little attempt has been made to speed up the optimization process, and it is possible that better methods could be developed in the future.

Besides retrieving the atmospheric profile, the posterior distribution may also be used to compute uncertainties in the retrieved parameters. There are a number of methods that could be used to obtain the uncertainties, depending on the definition of the error bars. Probably the most correct definition of the uncertainty comes from the marginal posterior distribution, that is, the one-dimensional function that remains after integrating over all the other parameters. The simplest method for computing this type of marginal error bar is to use the Monte Carlo integration method to find the second moment for each variable. If the maximum probability density function method is used, one may define a parameter uncertainty in which the other variables are held fixed at their retrieved values. If there is significant correlation between the parameters making up the precipitation structure profile, then this method would give smaller uncertainties than the more correct method. Because of the difficulties in interpreting the uncertainties with this method, error bars are not computed for the single-pixel retrievals done in this work.

3. Theoretical simulation test

The Bayesian precipitation retrieval method is a very flexible and complex system with choices to be made on the precipitation structure, the width of the forward probability distribution, and how the prior probability distribution is made. To determine the effect of these choices, theoretical tests are performed by simulating upwelling brightness temperatures from cloud model

fields and performing retrievals on these synthesized observations. The advantages of theoretical tests are that one may control the simulation and retrieval processes to look into particular issues and that the "true" hydrometeor profiles are known.

The cloud model output used for the theoretical test is from the two-dimensional GATE tropical squall-line simulation by Tao and Simpson (1989) using the Goddard cumulus ensemble (GCE) model. The GCE model used the microphysical scheme described in Rutledge and Hobbs (1984), which has two categories of liquid water (cloud water and rain) and three categories of ice phase (cloud ice, snow, and graupel). The precipitating hydrometeors have exponential distributions with fixed intercepts, so each species is specified by a single parameter, the mass content. Three regions from two model times (2 and 8 h) comprising 300 pixels are selected for the theoretical tests. The microwave radiative transfer simulations are carried out for 10.7, 19.4, 37.0, and 85.5 GHz for horizontal polarization at an observation angle of 53°. The full phase function information is computed from Mie theory. The radiative transfer is computed using the spherical harmonic spatial grid model (Evans 1993) in plane-parallel mode for land (emissivity of 0.9) and flat water (Fresnel) surfaces. To make the synthetic observations, 1-K rms Gaussian noise is added.

The purpose of the precipitation structure test is to determine how the accuracy of the retrieval depends on the number of layers and complexity of the hydrometeor representation. A precipitation structure provides the framework for the retrieval, specifying the thickness of the layers and which parameters are allowed to vary in the retrieval. The 10 structures (listed in Table 1) range from a two-layer ice-above-rain structure to a structure with variable rain, snow, graupel, cloud water, water vapor, and temperature. The retrievals are done with the forward probability distribution width σ_m set to 3 K for each channel. The prior probability distribution parameters are computed from 15 1-h cloud model intervals, using an integrated rain cutoff of 0.04 kg m^{-2} (about 1 mm h^{-1}).

The basic results of this test are rms differences in hydrometeor mass content between the retrieved profiles and the original cloud model. Figure 1 shows the error in surface rain mass for land and water retrievals as compared to the error of assuming climatology from the cloud model priors. As expected, the error in rain retrieval is much worse over land than water, except for area A, from the developing stage, which has small amounts of ice. The error over land is much worse for those structures that have no cloud water (1 and 4). The error is the same for structures 5 and above except for structure 9, which has increased error because of the greater variability of its 1-km-thick lowest layer. The retrieved rms error is a smaller fraction of the climatology error when measured logarithmically. The structures 5 and above are equivalent in the retrieved average integrated rain. The retrieval of integrated ice (not shown) is significantly more accurate for structures that include cloud water, and structures 5 and above have equivalent error. The fractional error for cloud water is much larger than for rain or ice, though the retrieval is considerably better than assuming a fixed cloud water amount.

The results of this test show that there is little to be gained from using precipitation structures more complex than the five-layer structure 5 with rain, ice, and cloud water. Simpler two- or three-layer structures do not appear to work as well. Comparison of structures 4 and 5 indicates that including cloud water is important. The extra variables (water vapor and temperature) of structure 10 neither help nor hurt, but it does take much longer to optimize with that many variables. These conclusions give some indication of what precipitation structures to use for Bayesian precipitation retrieval but cannot be generalized automatically to all other situations.

Tests of the other input choices in the retrieval method were performed in a similar manner. The re-

TABLE 1. Variable parameters in precipitation structure tests. The number of layers and variables is listed for each structure. The range (km) that each hydrometeor is variable is given ("N" means none, "C" means climatological value).

Precipitation structure	N_{lay}	N_{var}	Height range that parameter is variable					
			Rain	Snow	Graupel	Cloud	Vapor	Temperature
1	2	2	0–4	N	4–8	N	C	C
2	2	4	0–4	N	4–8	0–8	C	C
3	3	7	0–5	N	4–8	0–8	C	C
4	5	6	0–5	N	4–10	N	C	C
5	5	11	0–5	N	4–10	0–10	C	C
6	5	14	0–5	4–10	4–10	0–10	C	C
7	7	14	0–5	N	4–12	0–10	C	C
8	7	18	0–5	4–12	4–12	0–10	C	C
9	11	22	0–5	N	4–12	0–10	C	C
10	6	27	0–5	4–11	4–11	0–11	0–7	0–11

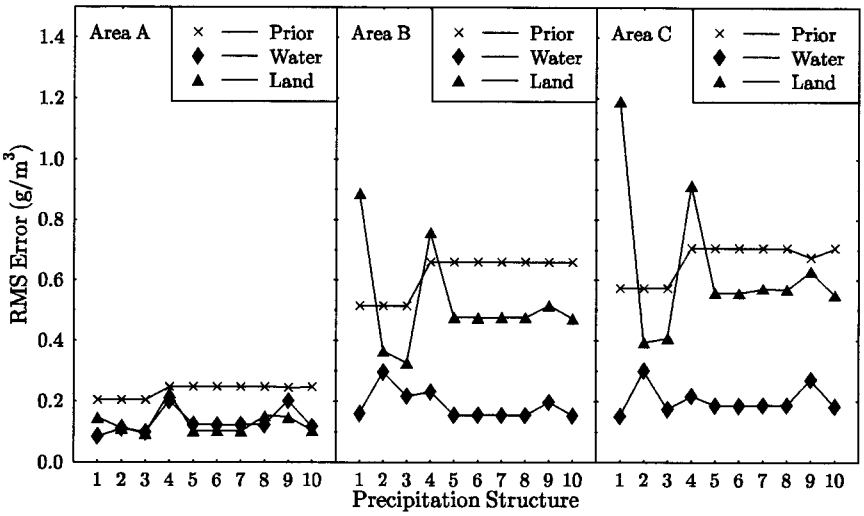

FIG. 1. Arithmetic rms retrieval error of surface rain mass content for the 10 precipitation structures for each of the three test areas. The retrieval error over water and land is compared with the error of assuming a fixed climatological value from the prior information.

trievals are not highly sensitive to the width of the forward probability distribution. The error increases if the observation width is much below the typical size of the forward modeling error and also increases for large widths as the retrievals move toward the climatological prior value. The covariance information in the prior probability distribution is important, but the retrievals are not very sensitive to the stage of storm development in the cloud model. These tests have a number of limitations, including the simplifications of the radiative transfer modeling and using the same cloud model output to make both the prior distribution and to synthesize the brightness temperatures. Real data will give worse retrieval results because of larger modeling errors and less applicable priors than was the case for these tests.

4. CaPE experiment validation test

Any proposed remote sensing inversion technique must be validated by using the method on field observations and comparing the results with other accepted measurement techniques. Precipitation remote sensing validation is somewhat problematic because there are no highly accurate area-averaged measurement techniques. Literal ground truth in the form of rain gauges offer only point measurements, and the very dense gauge networks needed for measurements in a collocated case study are impractical. Radar observations give the needed area and/or volume coverage but with uncertainties in rainfall estimation of perhaps a factor of 2 from uncertainties in drop size and precipitation inhomogeneities (Atlas and Ulbrich 1990; Joss and Waldvogel 1990). Polarimetric radars provide higher accuracy for rainfall estimation than radars that measure reflectivity only.

A major component of a space-based precipitation retrieval method that is not yet implemented in the Bayesian precipitation retrieval system is a way of dealing with beam filling, that is, the effects of variability within the radiometer field of view. Fortunately, the beam filling problem can be reduced substantially by using data from high-resolution aircraft-based microwave radiometers. Another reason for using high-resolution aircraft data is that the prior information for the retrieval is obtained from a limited amount of cloud model output with 1-km resolution.

The passive microwave data used for this validation are from the AMPR, which is the latest National Aeronautics and Space Administration (NASA) aircraft-based microwave instrument appropriate for precipitation sensing. The AMPR has a wider range of frequencies (from 10 to 85 GHz) than previous instruments, such as the Microwave Precipitation Radiometer, and so is better for testing multichannel algorithms. The AMPR has very high resolution (from 0.6 to 2.8 km), which should resolve much of the spatial variability of convective precipitation. The first organized field experiment that the AMPR flew in was the CaPE experiment in central Florida during the summer of 1991 (Williams et al. 1992). The AMPR and CP-2 data used for the validation test are from this experiment, during which collocated AMPR and CP-2 data of convective precipitation were acquired over the ocean and land.

a. The AMPR instrument and data

The AMPR has been designed and built in the last five years under the direction of the NASA Marshall

Space Flight Center (Spencer et al. 1994; see also Vivekanandan et al. 1993). The AMPR mounts in the ER-2 aircraft which cruises at 20 km, allowing it to overfly deep convection. The AMPR has the same multifrequency feedhorn as the Special Sensor Microwave/Imager (SSM/I) for its 19.35-, 37.1-, and 85.5-GHz channels and an additional specially designed feedhorn for 10.7 GHz. The 3-dB beamwidths are 8.0° for 10.7 and 19.35 GHz, 4.2° for 37.1 GHz, and 1.8° for 85.5 GHz, giving ground resolutions of 2.8, 1.5, and 0.6 km, respectively. The instrument scans across the track through nadir from +45° (right side) to −45° (left side), in 50 beam spots every 1.8°. The three lower frequencies are thus oversampled in that they are sampled multiple times per beamwidth. A new scan is started every 3.0 s, which at the 200 m s^{-1} cruising speed of the ER-2 corresponds to 0.6 km.

Cross-track scanning with a 45° offset flat-plate reflector causes the instrument's polarization basis to rotate relative to the scene orientation. While the single polarization is always linear it rotates from horizontal at $\theta = +45°$ (start of scan) to vertical at $\theta = -45°$ (end of scan). The polarization measured by AMPR thus varies with angle but is a linear combination of the vertical and horizontal polarization of the scene given by

$$T_b^{\mathrm{AMPR}} = T_b^H \cos^2(\theta - 45°) + T_b^V \sin^2(\theta - 45°). \quad (5)$$

The 50-ms spot integration time results in a receiver noise rms ranging from 0.15 to 0.30 K. After every fourth scan, the reflector rotates to view hot and cold calibration targets during one scan period. The absolute accuracy is estimated to be no better than several degrees, especially for the coldest brightness temperatures that must be extrapolated beyond the calibration temperature.

The CP-2 multiparameter radar is operated by the National Center for Atmospheric Research. It provides dual-frequency operation at S band (3 GHz) and X band (10 GHz) with matched 0.9° beamwidths (Bringi and Hendry 1990). Pulse-to-pulse polarization switching at S band provides for the measurement of differential reflectivity $Z_{\mathrm{DR}} = 10 \log(Z_{HH}/Z_{VV})$ in addition to the standard horizontal reflectivity Z_{HH}. At X band, the linear depolarization ratio LDR and X-band reflectivity can be measured. The X-band specific attenuation can be derived by comparing S- and X-band reflectivity.

Two AMPR datasets are used for the precipitation algorithm validation, one over land and one primarily over ocean. Uncooperative weather, the requirement that the precipitating systems be within CP-2 range, and several AMPR equipment failures severely limited the number of useful CaPE AMPR CP-2 datasets available for algorithm validation. Figure 2 shows a map of the experiment area with the location of the CP-2 radar, various upper-air stations, and the two

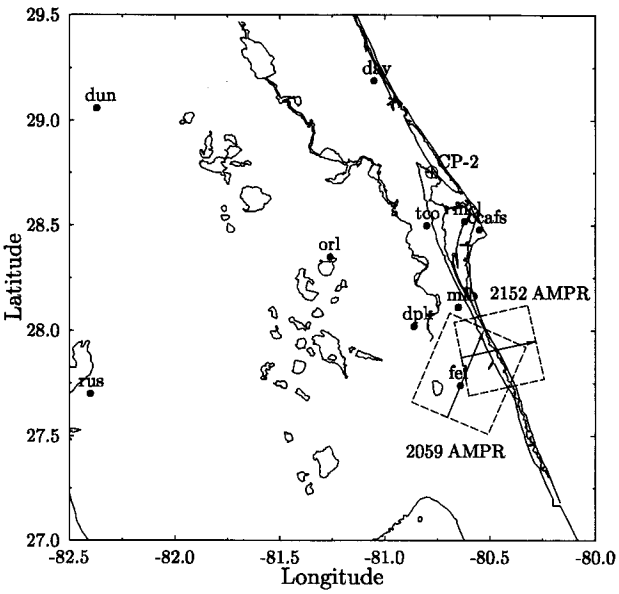

FIG. 2. Map of the CaPE experiment area with locations of the two AMPR datasets. The arrows in the dataset boxes show the heading of the ER-2. The CP-2 radar position and the locations of the upper-air stations are also shown.

AMPR datasets. The datasets were both from 12 August 1991 and are identified here by their starting times: 2059 UTC for the one over land and 2152 for the one over the ocean. The centers of the AMPR datasets are about 100 km from CP-2. There is an additional dataset available 40 km northwest of CP-2 (2226 UTC), which is not used here because the lower frequency channels are contaminated by the combination of land and ocean background. The two datasets are made by combining the AMPR and CP-2 data (Vivekanandan et al. 1993). The CP-2 data are remapped into the AMPR scanning coordinate system by tracing AMPR beams through a three-dimensional grid of CP-2 data with 0.5-km spacing and extracting the closest grid point for each 0.5-km altitude. The natural coordinate system for AMPR is the pixel-scan system. The pixel number from 1 to 50 (right side to left side) is the cross-track coordinate, while the scan number (starting at 1 in the dataset) is the along-track coordinate (in the direction the aircraft flies). Since the aircraft flew directly over the convective cells, the nadir pixel (25) is used in the following retrievals. Scans 1–60 of the 2059 UTC dataset and scans 30–60 from the 2152 UTC dataset are chosen to avoid the coastal scans.

Figure 3 shows the AMPR brightness temperatures along the nadir track for the ocean and land datasets. The ocean (2152 UTC) dataset has a precipitating cell near scan 35 with raised 10-GHz brightness temperatures and depressed 37- and 85-GHz temperatures. The 19-GHz channel is saturated near 250 K over the heavy rain area before scan 40. The elevated 19- and 37-GHz

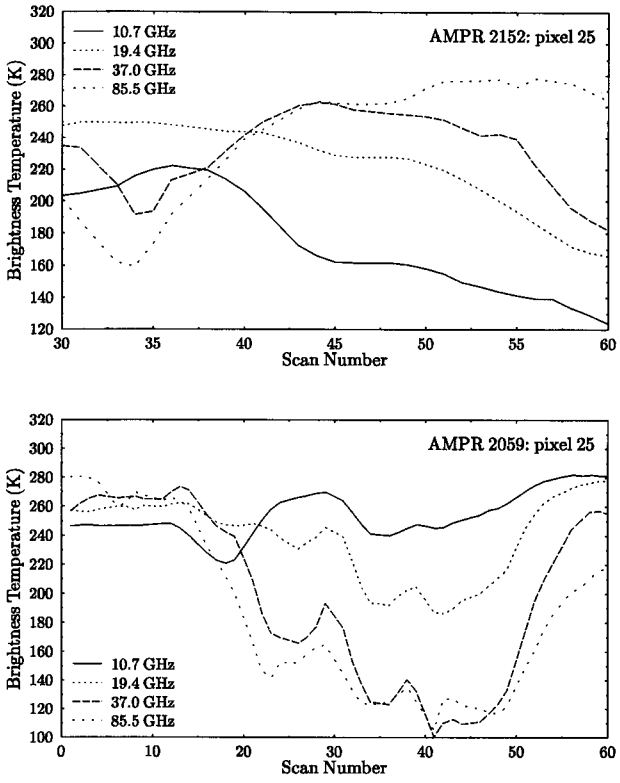

FIG. 3. AMPR brightness temperature traces for pixel 25 (nadir) in the 2152 UTC (ocean) and 2059 UTC (land) dataset. The scan direction is along the aircraft track.

brightness temperatures indicate lighter rain from scan 45 to beyond scan 55. Water vapor emission prevents the 85-GHz channel from going much below 270 K in the nonraining area. In the 2059 UTC dataset, the large brightness temperature depressions for the 37- and 85-GHz channels (to 100 K) show that the convection over land is much stronger than in the ocean dataset. There is also a substantial brightness temperature decrease at 19 GHz. The dip in the 10-GHz channel near scan 20 is probably due to wet ground. Further description of these two AMPR datasets and the collocated CP-2 radar data is available in Turk et al. (1994).

b. RAMS cloud model simulation

The prior probability information for the validation experiment is obtained from a simulation performed with the Colorado State University Regional Atmospheric Modeling System (RAMS) (Tripoli and Cotton 1982; Cotton et al. 1982; Cotton et al. 1986). The simulation was a modification of a two-dimensional summertime Florida sea-breeze simulation performed by Nicholls et al. (1991). The horizontal domain of 400 km has a land surface for the center 200 km. The horizontal resolution is 1 km, and vertical resolution ranges from 400 to 1000 m. The simulation was started at 0800 local standard time (LST) and run for 14 h. The simulation performed was like experiment 3 in Nicholls et al. (1991). Their type 3 wind profile was used for initialization as it was the most similar to the winds obtained from the CaPE upper-air stations on the morning of 12 August 1991. The initial profiles of temperature and water vapor were obtained from the upper-air soundings and differ from that of Nicholls et al. (1991), mainly in having less water vapor above 3 km.

The RAMS simulation used five categories of hydrometeors: liquid cloud droplets, rain, pristine ice crystals, aggregates, and graupel. The cloud droplets and pristine ice crystals have a monodisperse size distribution whose size varies. An exponential size distribution is assumed for the large hydrometeors, and only the mixing ratio is predicted. Another difference between the simulation described in Nicholls et al. (1991) and the one used here are the microphysical assumptions. In this simulation the average particle diameters were specified at 0.54 mm for rain, 3.30 mm for aggregate, and 1.2 mm for graupel. The bulk densities of the ice particles are 0.9 g cm^{-3} for graupel and 0.03 g cm^{-3} for aggregate. These values are the default RAMS microphysics, except for the somewhat larger graupel (1.2 mm vs 1.0 mm). The high density and large diameter of the "graupel" category is more like that of hail.

The sea-breeze circulation develops a precipitating cloud about 3 h into the simulation. The temperature, water vapor mass content, and hydrometeor mass contents from the simulation are stored every 15 min from 180 to 840 min simulated time. Figure 4 shows the total precipitating hydrometeors for six cloud model times at 2-h intervals. At first, the west coast sea-breeze front is stronger and produces a line of precipitation (240 min). By 360 min, the west coast precipitation has died out and the east coast front convection is producing rain. The convection moves west and rapidly gathers strength, perhaps from the collision of the two sea-breeze fronts (480 min). The storm then grows and develops a stratiform region, while the convective region is reduced in intensity (600 min). In the early evening, the system dissipates but still produces light rain.

c. Radar validation methods

There are two approaches to validating hydrometeor profiles retrieved from passive microwave measurements using radar observations. One is to retrieve profiles from the radar data and compare hydrometeor mass contents, for example. The other is to simulate radar observables from the retrieved size distributions and compare these. The advantage of the latter ap-

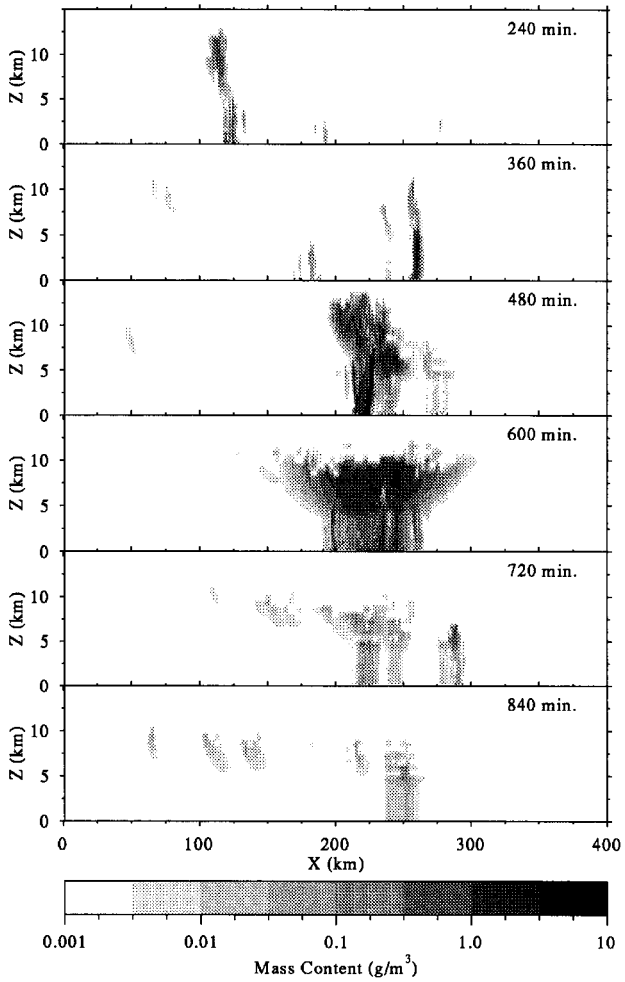

FIG. 4. Total precipitating hydrometeor (aggregate, graupel, rain) mass content for six model times from the Florida sea-breeze simulation after Nicholls et al. (1991). The mass content is displayed logarithmically from 10^{-3} to 10^{1} g m^{-3}.

proach is that the uncertainties of radar inversion are avoided. The disadvantage is that comparison is done in terms of radar parameters (e.g., reflectivity) instead of the desired end product of hydrometeor mass content or rainfall rate. Both approaches are used for this validation experiment.

The radar simulation computes reflectivity (dBZ) at S band. Attenuation at X band is not available for the validation, because the large radar range ($\approx$100 km) made the attenuation unreliable. The same microphysical assumptions as for the Bayesian retrieval method are used to derive the hydrometeor-size distribution and other properties. The Mie calculation and integration procedure is the same, but backscattering is computed and converted to reflectivity. Differential reflectivity Z_{DR} is not simulated, because the calculations assume spherical particles. The simulated reflectivity is compared with the average of the horizontal and vertical reflectivity, which is fairly close to the equivalent volume sphere reflectivity for oblate raindrops.

The process of inverting radar observables to find hydrometeor distributions is not unique and is an area of active research. Since S-band radar is sensitive only to larger sizes, cloud liquid water content cannot be obtained. Two radar inversion methods are used here. The method recently developed by Turk et al. (1994) retrieves the three parameters (N_0, Λ, and m) for a gamma size distribution $N(D) = N_0 D^m \exp(-\Lambda D)$ of equivalent volume spheres. The type of particle (water, ice, or melting) is assigned based on Z_{DR}, and a particle density of $\rho = 0.75$ g cm^{-3} is assumed for ice. The procedure starts near the surface where an N_0 and m are assumed and Λ computed to agree with the reflectivity. The method then moves upward and adjusts the three parameters smoothly, keeping agreement with the reflectivity. As the reflectivity decreases, Λ increases and m decreases. The profiles are processed starting from the core region and moving radially outward. The particle mass content is computed from the size distribution and compared with the passive microwave-derived mass content.

The other method retrieves both parameters of the exponential distribution but only for rain. Using the fact that the oblateness of raindrops is related to their volume, Seliga and Bringi (1976) showed how differential reflectivity could be used to infer the mean size of an exponential distribution of raindrops. With improved measurements of drop oblateness, Illingworth and Caylor (1989) computed simple fits to the D_0–Z_{DR} relation (D_0 is median volume diameter), which agreed closely with radar observations they presented. Table 2 lists the fit coefficients for D_0–Z_{DR} and also for Z_H–Z_{DR} with a fixed distribution intercept N_0. The fits give Z_H correc to 0.2 dBZ. The retrieval procedure computes the mean drop diameter and distribution intercept from Z_H and Z_{DR} with the formulas in Table 2 and then calculates the corresponding mass content. The major problem with this inversion method appears to be contamination of the rain signal with ice below the freezing level. Large, nearly spherical ice particles decrease Z_{DR} while increasing Z_H. This is interpreted as small average raindrop size with high reflectivity, which leads to extremely large computed mass contents. This problem is dealt with by limiting the exponential distribution intercept N_0 to a maximum value of 8000 mm^{-1} m^{-3}. In addition, the rain mass is retrieved only in areas below 4 km and with $Z_{DR} \geq 1.0$ dB.

d. Bayesian precipitation retrieval setup

The previously described procedures are used to perform the hydrometeor profile retrievals from the cloud model output and AMPR brightness tempera-

TABLE 2. Fit coefficients for D_0–Z_{DR} and Z_H–Z_{DR} relations for raindrops from Illingworth and Caylor (1989) (Table 3). The first fit (3.67 $\bar{D} = D_0 = \Sigma_i a_i Z_{DR}^i$; Z_{DR} in decibels, D_0 in millimeters) is used to find the exponential distribution slope or mean size $\bar{D}$ from Z_{DR}. The second ($Z_H = \Sigma_i a_i Z_{DR}^i$, for $N_0 = 8000$ mm^{-1} m^{-3}; Z_{DR} in decibels, Z_H in dBZ) is used to find the intercept N_0 by scaling N_0 to match Z_H. The fits assume a maximum drop diameter of 10 mm.

a_0	a_1	a_2	a_3	a_4	Z_{DR} range (dB)
0.4453	1.311	−0.9074	0.3863		0.1–1.0
0.5998	0.6762	−0.04640	0.003804		1.0–4.5
2.620	95.14	−162.8	159.0	−59.15	0.1–1.0
16.58	22.64	−5.020	0.6882	−0.03818	1.0–4.5

tures. A seven-layer precipitation structure with variable parameters listed in Table 3 is used. Only the hydrometeor mass contents are variable; the mean particle size is fixed. For the land case, the surface emissivity is variable (i.e., retrieved) and a prior distribution with mean of 0.85 and standard deviation of 0.05 is assumed. For the ocean case, pure Fresnel reflection is assumed. The ocean surface temperature is set to 301 K, while the land surface is at 305 K. The temperature and water vapor profiles are held fixed at their climatological values (from the cloud model).

It would seem to be desirable to have two categories of ice in the precipitation structure as the cloud model did. One category would be *graupel* or moderate-size high-density ice, while the other would be *aggregate* with larger-size low-density ice. Including the aggregate category (for example, with mean diameter 2.50 mm and density 0.1 g cm^{-3}) causes unstable retrievals because the observations cannot be matched reasonably. Forward radiative transfer modeling from the RAMS cloud model hydrometeor fields shows that the low brightness temperatures (<110 K for both 37 and 85 GHz) of the strong convection land case cannot be met. This is because nearly all the ice mass in the RAMS fields is in the aggregate category rather than the graupel category, and low-density particles do not scatter enough. For this reason, a single ice category is used, and the masses of aggregate and graupel are combined for determining the prior distribution from the cloud model fields. This can be thought of as using the equivalent volume sphere method of approximating low-density ice with smaller solid ice. The microphysical consistency of the cloud model output may be violated by having a single ice category, but the brightness temperature observations are indicating that aspects of the microphysical output from the cloud model may not be correct.

TABLE 3. The variable parameters in the precipitation structure used in the CaPE validation.

Layer	Heights (km)	Variable parameters
1	0–2	rain
2	2–4	rain, ice, cloud
3	4–5	rain, ice, cloud
4	5–7	ice, cloud
5	7–9	ice, cloud
6	9–11	ice, cloud
7	11–13	ice

Three sets of microphysical assumptions are made for the retrievals. The first is that of the RAMS simulation (rain mean diameter of 0.54 mm, ice diameter of 1.20 mm), the second has larger raindrops and smaller ice (0.70 and 1.00 mm mean diameter, respectively), and the third uses the Rutledge and Hobbs fixed intercept (N_0) distributions. Since the ice particles are modeled as high density (0.9 g cm^{-2}), the unscaled Eddington approximation is used. The prior distribution is made from all 45 model output times of the simulation. The cutoff for raining pixels is 0.04 kg m^{-2} of rain, which gives 1680 raining columns over the whole cloud simulation. The somewhat high rain cutoff is justified by this application to convective rainfall. In the results shown below, the width of the forward probability distribution is $\sigma = 10$ K for all channels, which usually gave somewhat better results than a width of $\sigma = 5$ K.

e. Precipitation retrieval results and comparisons

A comparison of CP-2 radar reflectivity with reflectivity simulated from the hydrometeor profiles retrieved for the 2152 UTC (ocean) dataset is shown in Fig. 5. This retrieval used the "standard" setup that has a mean diameter of 0.70 mm for rain and 1.00 mm for ice. The basic structure of the simulated reflectivity is roughly correct in terms of the location of the reflectivity maximum and the falloff with height. The reflectivity of the rain in the first cell is about right. The second rain cell, which is quite distinct in the CP-2 image, is not apparent in the AMPR-derived radar image. This is to be expected, since the 10-GHz brightness temperature stays level between scans 45 and 50. The AMPR-derived reflectivity is much too high in the nonraining area beyond scan 55. This is partly due to the prior distribution, made only from raining pixels, that causes there to always be some rain. The 10-GHz channel is still decreasing for the last few scans and has not reached its lowest value. This may be due to slight

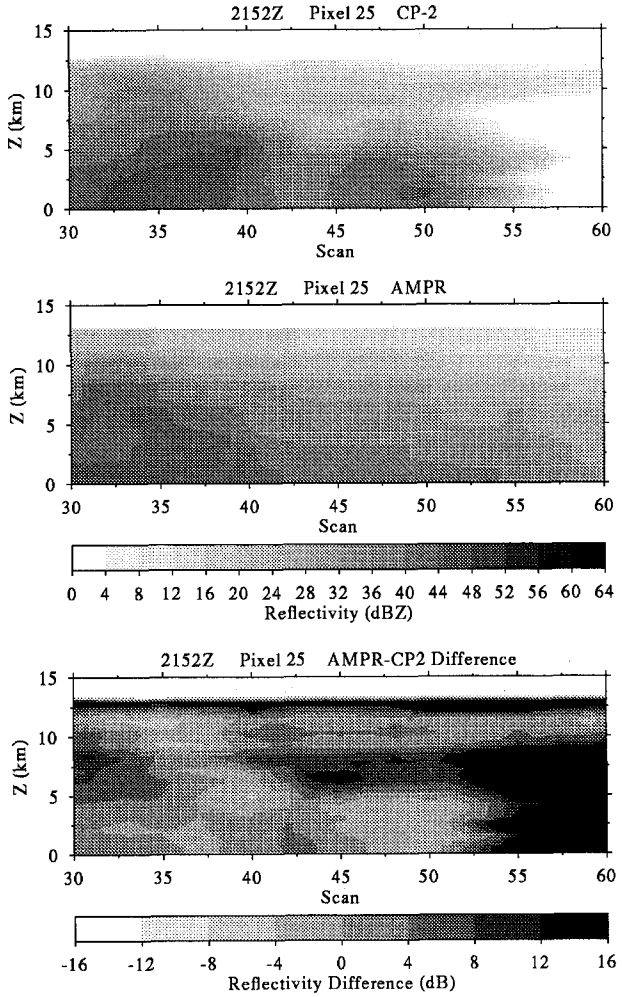

FIG. 5. Images of the reflectivity observed by CP-2 and reflectivity simulated from hydrometeor profiles retrieved from AMPR data for pixel 25 (nadir) in the 2152 UTC (ocean) dataset. The difference between the two reflectivity fields is also shown.

brightness temperature warming from cloud water, which the radar does not see, as well as the lower resolution of the 10-GHz channel. Not surprisingly, microwave radiometry cannot detect the detailed vertical structure if it does not follow the expectations of the prior distribution; for example, the retrieval misses the reflectivity minimum around scan 45 at a height of 8 km.

A more quantitative comparison of the observed and simulated reflectivity images is shown in the left panel of Fig. 6, which graphs the rms reflectivity difference profile. The rms is done for a subset of the data that has reflectivity above about 24 dBZ. In this water surface situation, the retrieval is relatively good, giving reflectivities in the rain layer to under 4 dB. Given the uncertainties in size distributions, retrieving rain reflectivity to a factor of 2 (3 dB) may be considered a success. The rms difference in reflectivity in the midlevels is higher (8–10 dB), mainly because of the overestimate of ice above the second cell.

A comparison of radar-derived and AMPR-retrieved vertically integrated hydrometeor mass content for the ocean dataset is shown in Fig. 7. The difference in rain mass between the two radar inversions is perhaps indicative of the error in the radar mass retrieval. The method of Turk et al. (1994) gives roughly equal integrated rain mass for the two rain cells ($\approx$2 kg m^{-2}), while using the formula of Illingworth and Caylor (1989) gives a much different rain mass for the two cells (4 vs 2 kg m^{-2}). These two inversions both match the reflectivity but use different raindrop distributions. The Z_{DR} is above 3 dB in the second cell, and the Illingworth and Caylor (1989) method retrieves mean raindrop diameters of 0.60–0.65 mm there, as compared to 0.50–0.55 mm in the first cell. The smaller raindrops in the first cell lead to a higher mass content in order to give the same reflectivity. The mean diameters of the Turk et al. (1994) method are larger and more nearly equal for the two cells.

The integrated mass content traces are shown for AMPR retrievals with the three microphysical assumptions discussed above. The fixed intercept assumption has the smallest mean diameter for raindrops and thus the largest mass contents, followed by the 0.54-mm and then the 0.70-mm mean drop diameters. The 10-GHz channel, which is providing most of the information for rain retrieved over the ocean, is clearly sensitive to the mean drop diameter and not only to the rain mass. The AMPR retrievals generally fall between the two radar estimates of integrated rain mass for the first precipitation cell. All of the microwave-retrieved rain masses are lower than the radar ones for the second cell. This is due to the lack of a second 10-GHz brightness temperature maximum, perhaps due to relatively fewer but larger drops causing an enhanced reflectivity without the associated brightness temperature warming. The passive microwave–retrieved integrated ice mass agrees relatively well with the radar-derived mass, except for the fixed intercept microphysics. The two retrievals with fixed mean particle size agree closely even though the average ice diameter is different by 20% (1.00 and 1.20 mm).

Turning now to the 2059 UTC dataset over land, Fig. 8 shows the reflectivity image comparison. The reflectivity is overestimated by the AMPR retrieval in the light precipitation area before scan 17. In the strong convection, the reflectivity in the lower ice layers is retrieved fairly well but is too low for the higher ice layers. This causes the retrieved reflectivity to look more stratiform, but this is an artifact caused by the prior information from the cloud model, as will be explained in the next section. The AMPR-derived reflectivity for rain in the two heavy precipitation cells is far too low compared with the actual CP-2 obser-

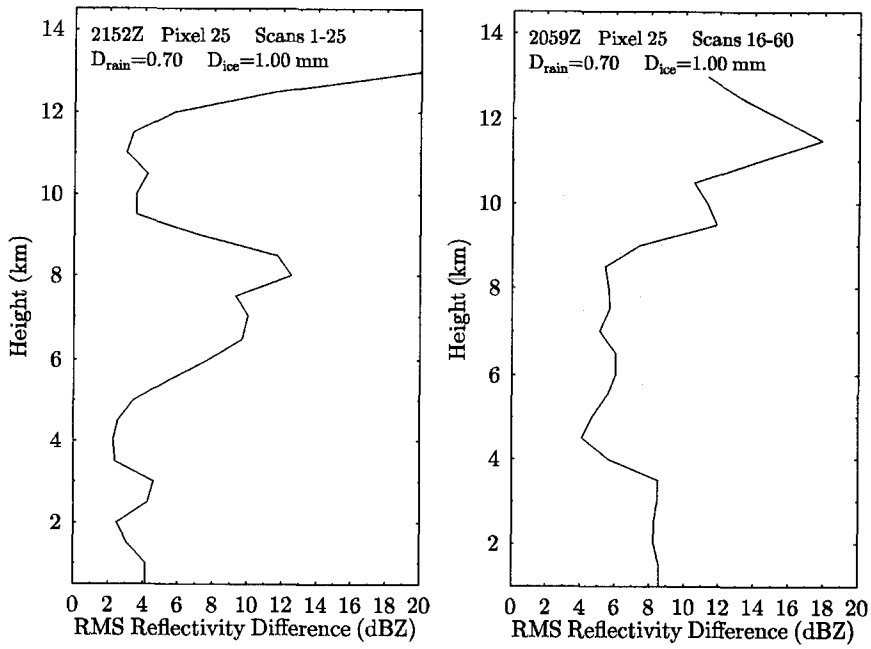

FIG. 6. Root-mean-square difference between the observed CP-2 reflectivity and the AMPR-derived reflectivity for the 2152 and 2059 UTC datasets.

vations. For retrievals over land, most of the information about the rain comes from its correlation with the ice mass through the prior probability distribution, because the rain is not sensed directly. Since the retrieval method operates on vertical columns, it also has difficulty with the slanted precipitation cells observed by the radar. The AMPR-derived reflectivity decreases only slightly in the area around pixel 30 that has a large drop in observed reflectivity. Over land, the change in 10-GHz brightness temperature from rain generally cannot be separated from surface emissivity changes, and here the retrieval of emissivity (not shown) closely follows the 10-GHz channel. For the land case, the retrieved reflectivity is further from observed than for the ocean case, as shown in the rms reflectivity difference profile (Fig. 6). The rms errors are around 8 dB in the rain layer, 6 dB in the lower ice levels, and much larger above.

Comparisons of CP-2 radar and AMPR microwave-derived integrated hydrometeor mass contents are shown in Fig. 9. All three microphysical assumptions for the microwave retrieval grossly underestimate the integrated rain mass. The Illingworth and Caylor (1989) radar inversion method for rain mass was heavily contaminated by ice (low Z_{DR} in high Z_H), and the correction technique leaves this estimate of rain mass fairly uncertain. The fixed-mean diameter AMPR results follow the radar-derived ice mass reasonably well, though falling short in the most intense cell. Even though variable cloud liquid water is in the precipitation structure, none is retrieved over land, although very small amounts (0.02 kg m^{-1}) are retrieved over the ocean (not shown).

5. Discussion of validation results and algorithm improvements

The AMPR retrievals made with the Bayesian method agree most closely with the radar results for the vertically integrated hydrometeors. For land and water background surfaces, the microwave-retrieved integrated ice mass content is fairly close to the radar-derived ice mass. The integrated rain mass content retrieved over the ocean is comparable to the radar-derived rain mass, but the amount of rain retrieved over land is generally much lower than the radar values. It should be remembered that there are significant uncertainties of perhaps a factor of 2 in the radar-derived hydrometeor mass, which limit the accuracy of the validation. The radar is, however, an independent source of detailed information about the hydrometeors. Other sources of error in the comparisons are the differing resolutions of the AMPR channels and the CP-2 radar and possible misnavigation of the AMPR data. The surface rain rate is not explicitly computed with a rain fallout model, because it is not directly relevant to the validation process.

The microwave retrieval of hydrometeor profiles as verified by the simulated radar reflectivity is much poorer than the integrated quantities. When the actual hydrometeor structure follows the profile expected from the prior information, the agreement can be good,

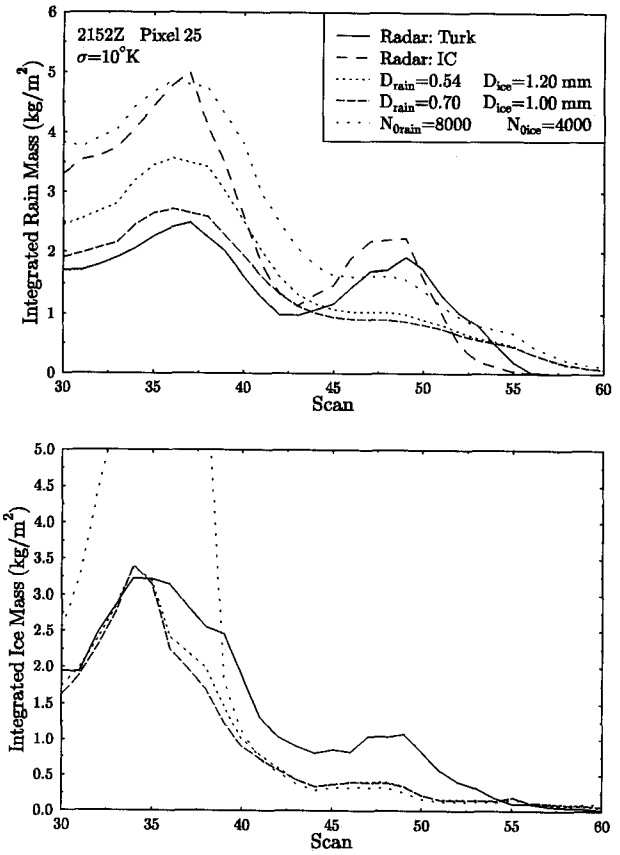

FIG. 7. Vertically integrated rain and ice mass content derived from CP-2 radar and retrieved from AMPR data pixel 25 in the 2152 UTC dataset. Three different microphysical assumptions made for the microwave retrieval are shown. The Turk radar method is that in Turk et al. (1994), and the IC method is from Illingworth and Caylor (1989).

but variations from the expected structure are not captured. This is to be expected from the physics of microwave radiative transfer through hydrometeors profiles. Microwave brightness temperatures are primarily sensitive to the integrated mass of ice or rain, depending on frequency, until the optical depth reaches the saturation level. Across the microwave spectrum, there are too few independent observations to retrieve a detailed vertical structure. Thus, much of the profile structure information must come from the prior distribution. The high sensitivity of the 10-GHz channel to surface properties means that over land there is little direct information about the rain mass in the microwave measurements.

For these reasons we do not expect the single-pixel retrievals of hydrometeor profiles to be particularly accurate. It still makes sense to retrieve profiles because various parts of the hydrometeor profiles are what the microwave radiation directly senses, and the vertical profile of hydrometeors is important for its connection to the latent heating profile. Sufficient accuracy in hydrometeor profiles and surface rainfall can come about only by averaging many single-pixel retrievals to find the climatological mean over a region. Of course, single-pixel retrievals can provide some useful information as to the likely rain intensity. But there are inherent limits to the accuracy of instantaneous microwave precipitation retrievals. This does not mean, however, that the retrieval results shown here cannot be improved upon significantly.

One advantage of the Bayesian framework is that one can examine explicitly the assumptions made in order to determine why the method behaved as it did and find ways to improve it. There are three parts to the Bayesian precipitation retrieval: the forward probability distribution, the prior probability distribution, and choosing the retrieved profile from the posterior distribution. The most important aspect to consider

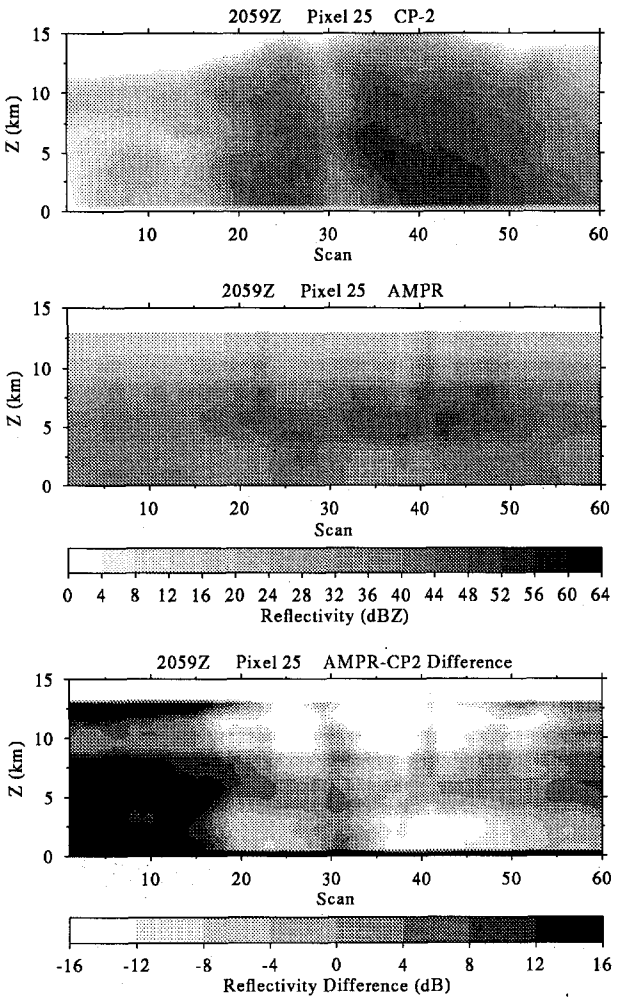

FIG. 8. Images of the reflectivity observed by CP-2 and reflectivity simulated from hydrometeor profiles retrieved from AMPR data for pixel 25 (nadir) in the 2059 UTC (land) dataset.

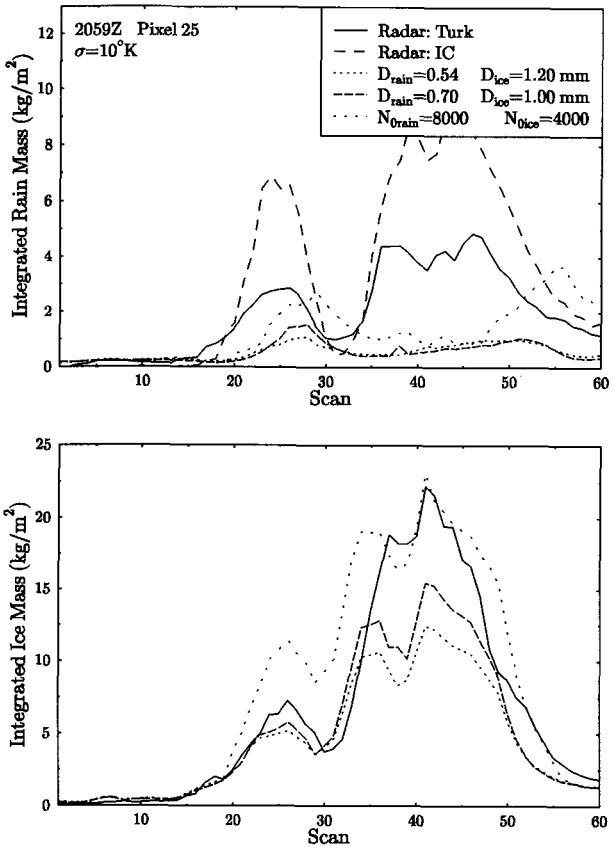

FIG. 9. Vertically integrated rain and ice mass content derived from CP-2 radar and retrieved from AMPR data pixel 25 in the 2059 UTC dataset. Three different microphysical assumptions made for the microwave retrieval are shown.

for the forward distribution are the errors in the radiative transfer modeling. For larger size parameters, there are uncertainties in modeling the scattering properties of low-density ice that could be investigated with more sophisticated scattering models. The assumptions made to derive particle size distributions from hydrometeor mass contents are the largest source of uncertainty. Of course, there will always be uncertainty in hydrometeor size distributions, but it would be desirable to have a cloud model predict two parameters of the size distributions to provide correct information about how distributions vary. With this prior information, both the mass content and the mean diameter of the hydrometeor size distributions could be variable in the retrieval process.

Another area of radiative transfer modeling error is from the surface emissivity. Here, very simple emissivity models are used: a single variable emissivity for land and a fixed Fresnel surface for water. The flat-ocean assumption underestimates the surface emissivity, so light rain may be retrieved in nonraining areas to compensate. The retrievals for land let the surface emissivity be variable over a large range, which probably reduces further the limited usefulness of the 10-GHz channel. The water emissivity model could be improved by using a model that depends on surface wind speed (e.g., Schluessel and Luthardt 1991), though there is considerable uncertainty in these types of models. The wind speed could then be retrieved, perhaps with the help of prior information connecting surface wind speed and area rain intensity. The land emissivity is much more difficult to model, but a model based on a variable wetness parameter, along with a known soil or vegetation type, might be possible. Having an additional low-frequency channel, such as the 6.8-GHz frequency on the Multifrequency Imaging Microwave Radiometer, may allow surface emissivity characteristics to be sensed directly.

The effects of horizontal variability on radiative transfer is another source of error. The plane-parallel model used here does not deal with three-dimensional variability, but the footprint filling problem is not likely to be too much of a problem for this AMPR retrieval because of the small footprints. In the future, the footprint filling problem could be dealt with by considering a distribution of precipitation intensities described by one or two parameters. These parameters could be related to the type or intensity of precipitation through three-dimensional cloud modeling, but an efficient way to perform the forward modeling would have to be developed. That the retrieval with the larger forward probability distribution width (σ = 10 K) agreed better in terms of simulated radar reflectivity implies that there are substantial forward modeling errors. In light of all the possible errors in the radiative transfer modeling, it does not make sense to fit the modeled brightness temperatures exactly to the observations.

Much of the reason for the errors in the profile retrievals can be explained by examining the cloud model output and prior distributions derived from it. There are two possible sources of error in the prior distributions. The first is the simplicity of the assumed functional form, that is, the multivariate lognormal distribution. The second is incorrectness or unrepresentativeness of the hydrometeor profiles produced by the cloud model. This issue is investigated by examining scatterplots of the 1680 hydrometeor profiles that went into making the prior distribution used in the retrievals. Figure 10 shows the relationship between the rain mass contents in the two lower layers and also between the lowest layer rain mass and the ice mass from 5 to 7 km. Along with the scatterplot, shown on the log–log plots, is the one-sigma ellipse derived from the covariance matrix that indicates the shape of the prior distribution. The two rain layers are well, but not completely, correlated. The correlation between the 0–2-km rain mass and the 5–7-km ice mass is quite low. The distribution of ice is clearly not lognormal, since it is not symmetric on this plot. The correlation be-

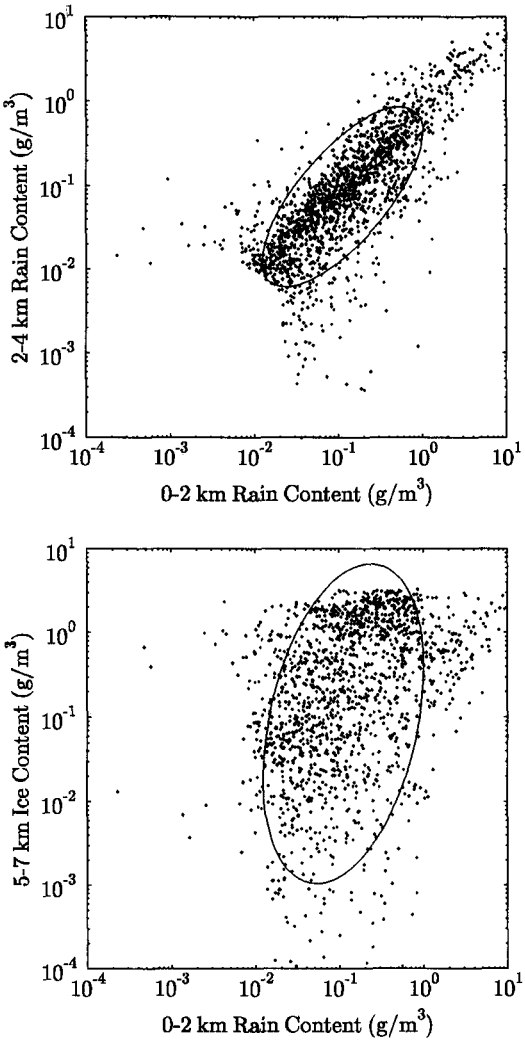

FIG. 10. Scatterplot of the hydrometeor mass contents produced by the RAMS simulation for various layers in the precipitation structure. Also shown is the one sigma ellipse indicating the lognormal prior distribution fit to the points. The top panel shows the rain mass for two layers, and the bottom panel shows the lowest layer rain and midlevel ice. The raining pixel cutoff eliminates the points in the lower-left corner of the top panel.

tween 2- and 4-km rain and cloud water is very low (Fig. 11). The cloud liquid water has a bimodal distribution that may be due to RAMS predicting it by determining what water mass is left after the conversion processes have acted. The result is that the log mean cloud water content is 0.002 g m^{-3}, which leads to very low retrieved cloud water contents. The 5–7-km ice-mass content, and that at 9–11 km, is moderately correlated, but the lognormal distribution is not representing the functional form of the scatter of points at all well. The scatterplot indicates that when there is a large amount of ice up high (strong updraft), then there is also much ice below, but there also can be a lot of ice in the lower region without much above. The underestimate of higher-level ice in the land retrieval is caused by this poor match between the form of the prior and the cloud model output. Figure 12 breaks the ice category down into the aggregate and graupel categories simulated in the RAMS run. As one would expect from microphysics, the correlation between rain and graupel is much higher than between rain and aggregate. However, the log-space mean of graupel mass is 70 times smaller than that of aggregate mass; this is why it was necessary to use the aggregate category. As a result, there is a low correlation between rain and ice, leading to the large underestimate of rainfall for the land retrieval.

Improvements in the microphysical parameterizations in cloud models used to generate the prior prob-

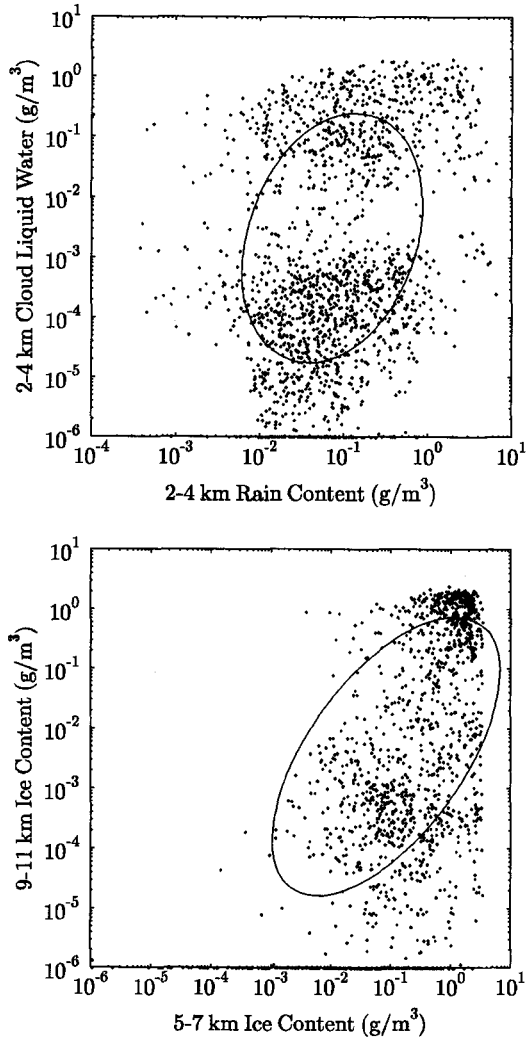

FIG. 11. Scatterplot of the hydrometeor mass contents produced by the RAMS simulation. The top panel shows the rain and cloud mass for one layer, and the bottom panel shows the ice mass for two layers.

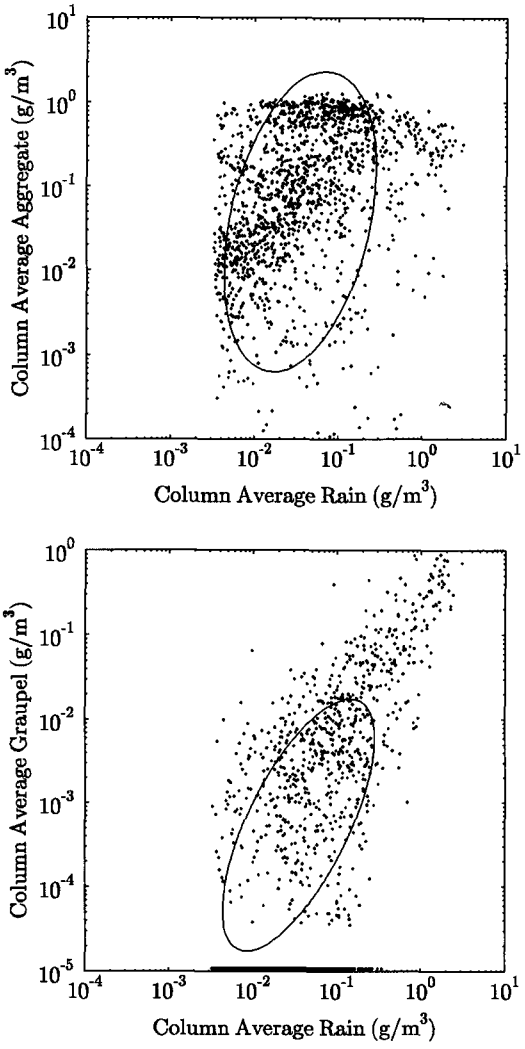

FIG. 12. Scatterplot of the hydrometeor mass contents produced by the RAMS simulation. The top panel shows the column-average rain and aggregate mass, and the bottom panel shows the column-average rain and graupel mass.

ability distribution should lead to improved hydrometeor retrievals with the Bayesian method. Bulk microphysical parameterizations that predict two parameters of the hydrometeor size distribution, and perhaps the ice density, are needed. Having a way to model ice- and water-mixture particle would be an additional improvement. The most important aspect of improved cloud modeling is that the microphysical parameterizations be tested for their effect on those aspects of hydrometeor distributions that are important for microwave radiation. Statistical validation of the parameterizations should be done using radar and passive microwave observations.

The algorithm validation with CaPE data performed here has shown some of the limitations of the simple multivariate lognormal form for the prior distribution. The lognormal distribution, while approximately correct, can lead to occasional large overestimates of quantities because of its long tail. The symmetry of the distribution in log space means that the very small values in the cloud model output can affect greatly the high end of the distribution through their influence on the distribution width. A more flexible way to describe the relationship between variables in the prior distribution could be used. Such a functional form will require more parameters and thus more cloud model data than the simple form used here. Close examination of cloud model output and measurements of hydrometeors could lead to a better functional form for the prior distribution.

Once the prior and forward distributions are determined, there is still the matter of how to choose the retrieved profile from the posterior distribution. The maximum of the probability density function depends in which space that function is measured. The current log-space density function may not be optimal and is fairly arbitrary. Other ways to use the posterior distribution, besides maximizing the density function, should be developed, especially for estimating area and time averages of precipitation. This validation effort should be considered preliminary because of the small amount of data and limited range of conditions considered. In addition to improvements to the existing algorithm, further work comparing hydrometeor profile retrievals with radar observations is needed.

6. Summary

A passive microwave precipitation remote sensing algorithm that retrieves hydrometeor profiles is described. Bayes theorem is used to combine a priori microphysical information from cloud models with forward radiative transfer modeling to select a physically likely profile that matches the observations. There is considerable flexibility as to the number of layers and the parameters that are retrieved in the hydrometeor profile. The prior probability distribution is assumed to have a multivariate lognormal form, which contains the correlations between hydrometeors at various levels that is important for a unique retrieval. The conditional probability distribution assumes independent, normally distributed brightness temperatures around those simulated with an Eddington radiative transfer model. A hydrometeor profile is retrieved by finding the maximum of the posterior probability density function computed from the prior and conditional distributions. The uncertainties in retrieved parameters also may be computed using the Bayesian framework.

A preliminary validation of the precipitation retrieval algorithm is carried out with data from the CaPE experiment in August 1991 in Florida. Hydrometeor profiles are retrieved over ocean and land from four

channel (10 to 85 GHz) observations from AMPR. Prior information is generated by a simulation of sea-breeze convection made with the RAMS cloud model. Validation is provided by data from the CP-2 multiparameter radar. Comparisons are made by simulating reflectivity from the AMPR-derived hydrometeor profiles and by inverting the radar data to mass content using two different methods. For these data over the ocean, the basic precipitation structure and amount is retrieved correctly, though too much precipitation is retrieved in the light and no-rain areas. Over land, the lower part of the ice layer and the integrated ice content is retrieved accurately, but too little rain is retrieved. The causes of the retrieval mistakes are explained in terms of the assumptions made in the algorithm and deficiencies in the cloud model microphysical output. Suggestions are made for improvements to the Bayesian precipitation retrieval algorithm, and the need to validate cloud model microphysical parameterizations is discussed.

Acknowledgments. This work was completed as part of a Ph.D. thesis by K. F. Evans under the supervision of G. L. Stephens. We acknowledge Ms. Robbie Hood and Dr. Roy Spencer at NASA Marshall Space Flight Center for providing the AMPR data. We thank Dr. Wei-Kuo Tao for providing the 2D GATE cloud model simulation output. Dr. M. Nicholls provided the setup for his RAMS Florida sea-breeze simulation. Dr. Hans Liebe provided his MPM92 code for computing microwave atmospheric absorption. Dr. Tim Cornwell of the National Radio Astronomy Observatory introduced one of us (Evans) to Bayesian inversion many years ago. Financial support for this research was provided by the National Aeronautics and Space Administration under Grant NAG-5-1592S. J. Turk acknowledges NASA Grant NAG8-890.

REFERENCES

Acton, F. S., 1990: *Numerical Methods that Work.* Mathematical Association of America, 549 pp.

Adler, R. F., H.-Y. M. Yeh, N. Prasad, W.-K. Tao, and J. Simpson, 1991: Microwave simulations of a tropical rainfall system with a three-dimensional cloud model. *J. Appl. Meteor.*, **30,** 924–953.

Atlas, D., and C. W. Ulbrich, 1990: Early foundations of the measurement of rainfall by radar. *Radar in Meteorology,* D. Atlas, Ed., Amer. Meteor. Soc., 86–97.

Bringi, V. N., and A. Hendry, 1990: Technology of polarization diversity radars for meteorology. *Radar in Meteorology,* D. Atlas, Ed., Amer. Meteor. Soc., 153–190.

Cotton, W. R., M. A. Stephens, T. Nehrkorn, and G. J. Tripoli, 1982: The Colorado State University three dimensional cloud/mesoscale model, 1982, II: An ice parameterization. *J. Rech. Atmos.*, **16,** 295–320.

——, G. J. Tripoli, and R. M. Rauber, 1986: Numerical simulation of the effects of varying ice crystal nucleation rates and aggregation processes on orographic snowfall. *J. Climate Appl. Meteor.*, **25,** 1658–1680.

Evans, K. F., 1993: Two-dimensional radiative transfer in cloudy atmospheres: The spherical harmonic spatial grid method. *J. Atmos. Sci.*, **50,** 3111–3124.

——, and G. L. Stephens, 1993: Microwave remote sensing algorithms for cirrus clouds and precipitation. Tech. Rep. 540, 198 pp. [Available from Colorado State University, Department of Atmospheric Science, Fort Collins, CO 80523.]

Fulton, R., and G. M. Heymsfield, 1991: Microphysical and radiative characteristics of convective clouds during COHMEX. *J. Appl. Meteor.*, **30,** 98–116.

Illingworth, A. J., and I. J. Caylor, 1989: Polarization radar estimates of raindrop size spectra and rainfall rates. *J. Atmos. Oceanic Technol.*, **6,** 939–949.

Joseph, J. H., W. J. Wiscombe, and J. A. Weinman, 1976: The delta-Eddington approximation for radiative flux transfer. *J. Atmos. Sci.*, **33,** 2452–2459.

Joss, J., and A. Waldvogel, 1990: Precipitation measurement and hydrology. *Radar in Meteorology,* D. Atlas, Ed., Amer. Meteor. Soc., 577–606.

Kedem, B., L. Chiu, and G. North, 1990: Estimation of mean rain rates: Application to satellite observations. *J. Geophys. Res.*, **95,** 1965–1972.

Kummerow, C. D., 1993: On the accuracy of the Eddington approximation for radiative transfer in the microwave frequencies. *J. Geophys. Res.*, **98,** 2757–2765.

——, and L. Giglio, 1994a: A passive microwave technique for estimating rainfall and vertical structure information from space. Part I: Algorithm description. *J. Appl. Meteor.*, **33,** 3–18.

——, and ——, 1994b: A passive microwave technique for estimating rainfall and vertical structure information from space. Part II: Applications to SSM/I data. *J. Appl. Meteor.*, **33,** 19–34.

——, R. A. Mack, and I. M. Hakkarinen, 1989: A self-consistency approach to improve microwave rainfall rate estimation from space. *J. Appl. Meteor.*, **28,** 869–884.

——, I. M. Hakkarinen, H. F. Pierce, and J. A. Weinman, 1991: Determination of precipitation profiles for airborne passive microwave radiometric measurements. *J. Atmos. Oceanic Technol.*, **8,** 148–159.

Liebe, H. J., 1989: MPM—An atmospheric millimeter wave propagation model. *Int. J. Infrared and Millimeter Waves,* **10,** 631–650.

——, G. A. Hufford, and M. G. Cotton, 1993: Propagation modeling of moist air and suspended water/ice particles at frequencies below 1000 GHz. *Proc. Atmospheric Propagation Effects through Natural and Man-Made Obscurants for Visible to MM-Wave Radiation,* AGARD Palma de Mallorca, Spain.

Liu, G., and J. A. Curry, 1992: Retrieval of precipitation from satellite microwave measurements using both emission and scattering. *J. Geophys. Res.*, **97,** 9959–9974.

Marzano, F. S., A. Mugnai, E. A. Smith, X. Xiang, J. Turk, and J. Vivekanandan, 1994: Active and passive microwave remote sensing of precipitating storms during CaPE. Part II: Intercomparison of precipitation retrievals over land from AMPR radiometer and CP-2 radar. *Meteor. Atmos. Phys.*, in press.

Mugnai, A., H. J. Cooper, E. A. Smith, and G. J. Tripoli, 1990: Simulation of microwave brightness temperature of an evolving hailstorm at SSM/I frequencies. *Bull. Amer. Meteor. Soc.*, **71,** 2–13.

——, E. A. Smith, and G. J. Tripoli, 1993: Foundations for statistical-physical precipitation retrieval from passive microwave satellite measurements. Part II: Emission-source and generalized weighting-function properties of a time-dependent cloud-radiation model. *J. Appl. Meteor.*, **32,** 17–39.

Nicholls, M. E., R. A. Pielke, and W. R. Cotton, 1991: A two-dimensional numerical investigation of the interaction between sea breezes and deep convection over the Florida peninsula. *Mon. Wea. Rev.*, **119,** 298–332.

Olson, W. S., 1989: Physical retrieval of rainfall rates over the ocean by multispectral microwave radiometry: Application to tropical cyclones. *J. Geophys. Res.*, **94,** 2267–2280.

Ray, P. S., 1972: Broadband complex refractive indices of ice and water. *Appl. Opt.*, **11,** 1836–1843.

Rutledge, S. A., and P. V. Hobbs, 1984: The mesoscale and microscale structure and organization of clouds and precipitation in midlatitude cyclones, XXI: A diagnostic modeling study of precipitation development in narrow cold-frontal rainbands. *J. Atmos. Sci.*, **41,** 2949–2972.

Schluessel, P., and H. Luthardt, 1991: Surface wind speeds over the North Sea from Special Sensor Microwave/Imager observations. *J. Geophys. Res.*, **96,** 4845–4853.

Seliga, T. A., and V. N. Bringi, 1976: Potential use of radar differential reflectivity measurements at orthogonal polarizations for measuring precipitation. *J. Appl. Meteor.*, **15,** 69–76.

Smith, E. A., A. Mugnai, H. J. Cooper, G. J. Tripoli, and X. Xiang, 1992: Foundations for statistical–physical precipitation retrieval from passive microwave satellite measurements. Part I: Brightness-temperature properties of a time-dependent cloud-radiation model. *J. Appl. Meteor.*, **31,** 506–531.

Spencer, R. W., R. E. Hood, F. LaFontaine, and E. A. Smith, 1994: High-resolution microwave imaging of the earth with the advanced microwave precipitation radiometer. Part I: System description and sample imagery. *J. Atmos. Oceanic Technol.*, submitted.

Tao, W.-K., and J. Simpson, 1989: Modeling study of a tropical squall-type convective line. *J. Atmos. Sci.*, **46,** 177–202.

——, ——, S. Lang, M. McCumber, R. Adler, and R. Penc, 1990: An algorithm to estimate the heating budget from vertical hydrometeor profiles. *J. Appl. Meteor.*, **29,** 1232–1244.

Tripoli, G. J., and W. R. Cotton, 1982: The Colorado State University three dimensional cloud/mesoscale model, 1982, I: General theoretical framework and sensitivity experiments. *J. Rech. Atmos.*, **16,** 185–220.

Turk, J., J. Vivekanandan, F. S. Marzano, R. E. Hood, R. W. Spencer, and F. J. LaFontaine, 1994: Active and passive microwave remote sensing of precipitating storms during CaPE. Part I: Advanced Microwave Precipitation Radiometer and polarimetric radar measurements and models. *Meteor. Atmos. Phys.*, in press.

Vivekanandan, J., J. Turk, G. L. Stephens, and V. N. Bringi, 1990: Microwave radiative transfer studies using combined multiparameter radar and radiometer measurements during COHMEX. *J. Appl. Meteor.*, **29,** 561–585.

——, ——, and V. N. Bringi, 1993: Comparisons of precipitation measurements by the Advanced Microwave Precipitation Radiometer and multiparameter radar. *IEEE Trans. Geosci. Remote Sens.*, **31,** 860–870.

Warren, S. G., 1984: Optical constants of ice from the ultraviolet to the microwave. *Appl. Opt.*, **23,** 1206–1225.

Weinman, J. A., and P. J. Guetter, 1977: Determination of rainfall distributions from microwave radiation measured by the *Nimbus 6* ESMR. *J. Appl. Meteor.*, **16,** 437–442.

——, and R. Davies, 1978: Thermal microwave radiances from horizontally finite clouds of hydrometeors. *J. Geophys. Res.*, **83,** 3099–3107.

Wilheit, T. T., A. T. C. Chang, M. S. V. Rao, E. B. Rodgers, and J. S. Theon, 1977: A satellite technique for quantitatively mapping rainfall rates over the oceans. *J. Appl. Meteor.*, **16,** 551–560.

Williams, S. F., K. Caesar, and K. Southwick, 1992: The Convective and Precipitation/Electrification (CaPE) experiment: Operations summary and data inventory. Tech. document, National Center for Atmospheric Research, Boulder, CO, 425 pp.

Yeh, H.-Y. M., N. Prasad, R. A. Mack, and R. F. Adler, 1990: Aircraft microwave observations and simulations of deep convection from 18 to 183 GHz. Part II: Model results. *J. Atmos. Ocean. Technol.*, **7,** 392–410.